

THE IN A HURRY SERIES

Behavioural Economics

in a Hurry

Humans aren't rational, profit from it

Max Brocklesby

BOOKSINAHURRY.COM

Contents

The Whole Thing in One Page 4

Why You Should Care 5

The Core Ideas 8

How It Actually Works 21

What People Get Wrong 30

Use It 35

Terms 40

Go Deeper 44

Notes and Sources 45

Bibliography 50

The Whole Thing in One Page

Behavioural economics is often introduced as a collection of tricks: anchoring, loss aversion, defaults, framing, overconfidence. That makes for good dinner conversation and bad economics. A named bias tells you little until you know the benchmark it violates, the mechanism producing it, the environment that sustains it and who gains from the resulting choice.

Standard economics begins by simplifying the chooser. Give a person preferences, beliefs and constraints, require enough consistency to rank the available options, and much of economic life becomes modelable. The simplification is powerful because it ignores attention, memory, temptation, social meaning and the way options are presented. Behavioural economics asks what happens when those omitted parts affect prices, contracts, saving, bargaining and markets.

The answer is not that humans are irrational in every direction. It is that choice is produced by a person in an environment. Attention is scarce, so people use shortcuts. Those rules can be efficient, yet they fail predictably when the cue they rely on is misleading. Outcomes are judged against reference points, so £100 gained and £100 not lost need not feel equivalent. Probabilities are distorted, frames move attention, and mental accounts make identical pounds behave differently. Defaults matter because doing nothing has consequences, and switching demands effort, confidence or notice.

Time divides the chooser too. The person who plans on Sunday can disagree with the person facing temptation on Friday. That creates demand for commitment devices, automatic saving and deadlines, while also creating revenue from forgotten renewals and delayed cancellation. Social preferences add another layer: people punish unfairness, reward reciprocity and care about intentions, so profit seeking is constrained by norms as well as prices.

These mechanisms become economically important when institutions scale them. A firm can make comparison easy and earn trust, or make

cancellation painful and earn an error rent. A government can simplify a benefit form, or use a fashionable nudge where a price, right or cash constraint is the real problem. Competition can eliminate exploitation when consumers learn quickly and rivals can profit from clarity. It can preserve or intensify exploitation when products are complex, feedback is slow and every seller benefits from confusion.

That is the subtitle's real promise. You can profit from behavioural economics in three ways: reduce the customer's decision cost, design commitment around goals the customer already has, or harvest predictable mistakes. The first two can create durable value. The third can work, but its durability depends on opacity, switching costs, weak feedback and regulation. Behavioural economics is therefore as much a theory of business models and institutions as a theory of minds.

The modern field learned to distrust its greatest hits. Some classic effects are robust; others are smaller, more conditional or more procedure-sensitive than popular accounts suggest. Published nudge effects have often exceeded effects in large operational trials. A behavioural claim is not a law of nature. It is a prediction about a mechanism in a specified setting.

The deepest problem comes last. Once observed choice can be changed by architecture, revealed preference can no longer settle welfare by itself. The designer must ask not merely whether behaviour moved, but whose objective defined success, whether the person would endorse the outcome under clearer conditions, and what happened after the first click.

That is the book.

Why You Should Care

A health-club contract is a compact experiment in economics, optimism and inertia.

Stefano DellaVigna and Ulrike Malmendier studied 7,752 members at three American gyms. Customers on monthly contracts costing more than \$70 attended about 4.3 times a month. Their expected cost per visit exceeded

\$17, while a ten-visit pass cost \$10 per visit. Across memberships, the average forgone saving was about \$600. The flexible monthly plan had another twist: its users were more likely than annual members to remain enrolled beyond a year, despite paying extra for the ability to cancel.

The members were not forced, deceived about the headline price or incapable of arithmetic. They bought for an active future self, then renewed as a less active present self. The transaction mixed a genuine desire to exercise with optimism about future attendance, procrastination and automatic continuation. If you read the purchase as a pure statement of stable preference, you miss the economic mechanism that made the contract profitable.

The same structure appears everywhere. A pension default can determine whether saving begins. The first number on a page can change the next number's meaning. A free trial can remove perceived risk at entry while transferring inertia to the seller at renewal. A one-page form can separate intention from completion. A programme described in lives saved can attract a different response from the identical programme described in deaths. Income and formal options may be unchanged, yet behaviour moves.

The commercial importance is larger than a collection of clever interface tricks. Every business chooses a menu, a comparison frame, a billing interval, a default, a cancellation path and a moment when the customer must pay attention. Those design choices determine how much of the firm's revenue comes from creating value and how much comes from customer error. A product that reduces search, uncertainty or self-control costs can earn behavioural profit even when customers fully understand it. A product that depends on forgotten renewal, hidden add-ons or exhausting exit earns an error rent that may last only while feedback, competition or regulation remains weak.

That distinction matters because competition does not automatically rescue the consumer. If one firm makes a complex product transparent while rivals hide profitable fees, transparency can win business. But if customers shop mainly on the visible headline price, every firm may have an incentive to

hide the rest. The market can then compete on the quality of the trap. Behavioural economics helps explain when learning and rivalry discipline exploitation and when they reproduce it.

Governments face the same architecture problem without the profit motive. Tax letters, benefit claims, organ-donation systems, pensions and school applications all have defaults and frictions. Small design changes can be cheap and useful, especially at scale. They can also distract from binding constraints. A reminder cannot make an unaffordable bill affordable. A default cannot replace a missing legal right. A simplified disclosure may be weaker than banning a hidden fee.

The personal lesson is therefore narrower and more useful than becoming a bias spotter. Ask what the benchmark is. Identify the reference point. Separate preference from friction. Notice when today's self and tomorrow's self have different incentives. Follow the designer's payoff. Then demand evidence that the mechanism works here rather than assuming a famous effect travels intact.

Behavioural economics becomes powerful when it stops saying people are irrational and starts explaining exactly how a predictable departure enters an economic system, with incentives attached to it. Once you can see that, contracts, prices, policies and your own choices look less like isolated decisions and more like engineered environments with incentives on both sides.

The Core Ideas

1. The Rational Actor Is a Benchmark, Not a Biography

The person in a standard economic model is stripped down for a reason. She has preferences, beliefs and constraints. Her preferences are complete enough to rank the relevant options and consistent enough that if she prefers A to B and B to C, she does not suddenly prefer C to A. She uses available information to form beliefs, then selects the feasible option that best serves what she wants.

Nothing in that description requires delight in spreadsheets. An economist can model a person as choosing rationally without claiming that she solved equations in the supermarket. The claim is often an as-if claim: her behaviour can be represented as though she maximised an objective. Nor does rationality mean selfishness. A person can care about family, strangers, principle or fairness and still choose coherently. The content of the preference and the consistency of the choice are different questions.

This abstraction earns its keep. If choices reveal stable preferences, economists can infer demand from purchases, predict responses to prices and compare policies without opening anyone's head. If a commuter keeps buying the slower but cheaper ticket, the model records a stronger preference for money than time at the relevant margin. It does not need a theory of childhood, mood or attention.

The trouble begins when the choice is not a clean reading of the preference. A person may misunderstand the probabilities, fail to notice a fee, postpone a form, copy a default or choose a contract for the person she expects to become. The observed action then mixes preference with belief, attention, friction and self-control. Treating it as pure preference can make a precise model of the wrong thing.

Herbert Simon attacked the assumption from another direction. Real decision-makers face limited time, knowledge and computational capacity.

They rarely inspect every option and calculate the best. They search until they find one that is good enough, a process he called satisficing. That can be intelligent. Searching has costs, and the perfect answer may be worth less than the time spent finding it. A rule that stops at a satisfactory hotel, supplier or job applicant can be rational once the cost of further search enters the model.

This is why behavioural economics does not replace rationality with chaos. It keeps the benchmark and studies patterned departures from it. The benchmark tells you what requires explanation. Without it, every choice is merely something that happened. With it, you can ask whether the gap arose from a mistaken belief, a limited search, a reference point, a social motive or an institution designed to make one action easier than another.

The benchmark also protects against lazy diagnosis. Calling a choice irrational requires more than disliking it. You must say which objective the person appears to hold, what information and constraints were available, and which consistency condition was violated. A poor person paying a high price for liquidity may be responding sensibly to a hard constraint. A rich person ignoring a trivial saving may be valuing attention. Behaviour becomes a puzzle only after the economic situation has been specified.

Core Idea 1 therefore sets the discipline of the whole book. The rational actor is not the enemy. It is the ruler held beside the choice. Behavioural economics begins when the distance from that ruler is systematic enough to model and consequential enough to change an economic outcome.

2. A Scarce Mind Uses Shortcuts

Suppose you had to price every risk by recalling all relevant evidence, weighting each observation by reliability and updating a probability distribution before deciding. You would still be comparing travel insurance when the flight took off.

A mind with limited attention must compress. Heuristics are the compression rules. They ignore some information, substitute an easier question for a harder one, or stop the search once a useful cue appears. The classic

research programme led by Amos Tversky and Daniel Kahneman concentrated on three families: representativeness, availability and anchoring.

Representativeness judges probability by resemblance. A new company with a charismatic founder, fast growth and a fashionable product looks like the story of a future giant, so the base rate of failure recedes. A job candidate who resembles the last successful hire feels safer than the statistics justify. Similarity contains information, but it does not contain all the information. When the underlying category is rare, even an impressive match may leave the event unlikely.

Availability judges frequency or risk by how readily examples come to mind. Recent layoffs can make recession feel imminent. A vivid product failure can dominate thousands of quiet successes. Memory is not a random sample of reality. It favours the recent, emotional, repeated and imaginable, which means the feeling of evidence can move without the evidence moving.

Anchoring begins from an available number and adjusts away from it. The first salary mentioned in a negotiation, the suggested donation on a form or the list price above a sale price can shift the eventual judgement. Some anchors carry information. Others should be irrelevant. Adjustment tends to be incomplete because constructing a new estimate takes effort and because the initial number changes which possibilities receive attention.

The easy conclusion is that shortcuts make people foolish. That conclusion is too easy. A heuristic can be well matched to an environment. Recognition can be useful when recognition tracks size or quality. Ignoring weak predictors can improve forecasts when data are noisy and samples are small. A simple rule can beat a complicated model that has fitted yesterday's accidents too closely. Gerd Gigerenzer and other critics of the biases programme made this correction forcefully: less information and less computation can sometimes produce more accuracy.

The economic question is therefore not whether a heuristic is rational in the abstract. It is whether the cue structure of the environment rewards it. A

shopper who buys the familiar brand may be saving search costs in a stable market. The same rule becomes exploitable when familiarity can be purchased through repetition while quality remains hidden. A lender who relies on a rough risk category may process applications cheaply. The rule becomes destructive when the category captures history rather than the applicant's capacity to repay.

This environmental fit explains why incentives do not erase every bias. Paying more for a correct answer can increase effort, but effort cannot recover information that was never noticed or a base rate that was never represented. Experience can help when feedback is prompt and clear. It helps less when outcomes arrive years later, noise masks the cause or each decision is unique. Property purchases, retirement choices and medical risks give the learner few clean repetitions.

Heuristics turn scarce attention into action. Their speed is the benefit; their selective blindness is the price. Behavioural economics becomes useful when it identifies which cue a decision rule is following, why that cue usually works, and what has changed in the environment to make it fail.

3. Value Starts From a Reference Point

Classical accounts of risky choice begin with final outcomes. If two gambles leave you with the same possible levels of wealth and the same probabilities, they should be valued in the same way. How the options were described or where you started should not matter.

Prospect theory changed the unit of experience. Daniel Kahneman and Amos Tversky argued that people commonly evaluate outcomes as gains and losses from a reference point rather than as final wealth. The reference might be the status quo, an expectation, a recent price, a target or what another person received. Move the reference and the psychological outcome can change while the pounds do not.

Imagine being told in January that your annual bonus will be £10,000. In December you receive £8,000. Your wealth has risen, yet the event is experienced partly as a £2,000 loss against expectation. Another worker

expecting nothing receives the same £8,000 as a gain. A model based only on final wealth misses the difference because the relevant comparison is carried inside the chooser.

Prospect theory adds two further shapes. Sensitivity diminishes with distance from the reference point. The difference between £0 and £100 usually feels larger than the difference between £10,000 and £10,100. The value function is also steeper for losses than for gains around the reference point, the pattern called loss aversion. Giving up an expected benefit can therefore matter more than receiving an equal extra benefit.

This does not mean every loss hurts exactly twice as much as every gain pleases. The famous ratio is a rough parameter associated with particular models and experiments, not a constant of human nature. Reference points are often uncertain, experience changes reactions, stakes and goods matter, and some procedures that appear to reveal loss aversion also contain misunderstanding or transaction friction. The durable claim is asymmetry and reference dependence in many settings, not a universal multiplier.

People also transform probabilities. Small probabilities can receive more decision weight than their numerical size suggests, while moderate and high probabilities may be compressed. That pattern helps explain why the same person can buy insurance against a rare loss and a lottery ticket for a rare gain. The person is not applying one stable attitude called risk aversion. The shape of the outcomes, their signs relative to the reference point and the probabilities all matter.

The change in model produces a striking reversal. People often prefer certainty to a gamble when choosing among gains, yet become willing to gamble when every option is framed as a loss. A sure loss closes the account painfully. A risky option offers a chance of escaping the loss, even if it also carries a worse outcome. This helps explain why failing projects attract another round of money and why investors may hold losing assets while selling winners, though any field application must identify the reference point rather than assume it.

Reference dependence matters wherever expectations can be managed. A pay rise below the expected rise can damage morale. A temporary discount can reset the price customers regard as normal. A guarantee can move the feared outcome from paying for failure to losing an assured benefit. Targets can motivate while they remain within reach and distort behaviour when falling short is coded as a loss to be avoided at any cost.

The central shift is simple to state and difficult to contain. Value does not sit inside an outcome waiting to be read. It emerges from a comparison. To understand a decision, identify what the person gets, then identify the result the person expected, owned, targeted or imagined instead.

4. The Same Options Can Produce Different Preferences

Rational choice requires invariance. If two descriptions lead to the same outcomes with the same probabilities, a person who understands both should not reverse preference merely because the words changed. Yet descriptions alter what receives attention and which reference point becomes natural.

The best-known demonstration concerned an imagined disease expected to kill 600 people. One pair of programmes was described in lives saved, another mathematically identical pair in deaths. In the gain frame, most participants preferred the certain option. In the loss frame, most preferred the gamble. Nothing in the outcomes had changed. The frame changed whether certainty looked like securing a gain or accepting a loss.

Framing is more than wording. It determines the mental account into which an outcome is entered. Richard Thaler used mental accounting to describe the informal bookkeeping by which people label money, combine or separate transactions, and decide when an account feels closed. A £100 theatre ticket lost on the way to the venue may deter replacement because the theatre account would now show £200. Losing £100 in cash on the same journey may not, although total wealth is identical. The category changes the perceived price of continuing.

These accounts can help. A household that separates rent, food and leisure money creates rough control over spending without calculating lifetime utility after every purchase. A business that assigns budgets can delegate decisions. The violation of fungibility is the mechanism that gives the rule force. Once money is labelled, moving it becomes psychologically costly even when the transfer is financially sensible.

Ownership can create another frame. In early endowment-effect experiments, people randomly given a mug often demanded more to sell it than unendowed people offered to buy it. Giving up the object was coded as a loss rather than declining to acquire it. Later work showed that procedures, understanding, trading experience and the type of good can narrow or remove the gap. The effect is not a magic property transferred by touch. It is evidence that entitlement, transaction meaning and reference points can enter valuation.

Defaults operate through a related set of mechanisms. An option selected in advance may become the status quo, an implied recommendation, the path requiring least effort or the choice that avoids responsibility for an active mistake. In a retirement plan, changing from opt-in to opt-out leaves the formal menu intact but moves the consequence of inaction from non-enrolment to enrolment. Persistence at the preselected contribution does not prove workers regard it as ideal. It may show that making a new choice is costly in attention or confidence.

This is where revealed preference becomes delicate. If a person chooses differently after a frame, account or default changes, which action reveals the true preference? The first answer cannot be 'the choice', because there are now several. Behavioural economics must distinguish preferences from mistakes without assuming that the researcher always knows better.

Sometimes the criterion is straightforward. A hidden fee that the buyer did not notice does not reveal a taste for paying it. A form abandoned because a button failed does not reveal disinterest. Harder cases involve genuine trade-offs among convenience, commitment, identity and freedom. The default may help a person carry out a stable goal, or it may install the

designer's goal in its place.

The lesson is not that preferences are unreal. It is that they are often incomplete until a situation asks for a choice. The menu, description, sequence and status quo help construct the answer. Economics can no longer treat those features as neutral packaging, because the package participates in the preference it appears merely to reveal.

5. Your Future Self Is Another Economic Actor

Standard models of choice across time use exponential discounting. A reward loses value at a constant proportional rate as it moves into the future. If you prefer £110 in eight days to £100 in seven days, moving both rewards forward by a year should not reverse the ranking. The distance between them remains one day.

People often behave otherwise when immediacy enters. £100 now can defeat £110 next week even when £110 in fifty-three weeks defeats £100 in fifty-two. The extra attraction belongs to now, not to the one-week interval. Models of present bias give the current period a special weight, producing preferences that can change as time passes.

That creates an internal bargaining problem. On Sunday, the planning self wants exercise on Tuesday. On Tuesday evening, the acting self wants the sofa. At the start of the month, the subscriber intends to cancel an unused service before renewal. On the final day, the task feels minor enough to postpone. Each decision is understandable in its moment. Together they defeat the plan.

The health-club data make the conflict visible because the contract records optimism about a future self and attendance records the future self's response. Members paid for frequent exercise, attended infrequently and delayed cancellation. A standard model can explain a high flat fee if people value flexibility or expect heavy use. It struggles when expectations, attendance and cancellation all point towards overconfidence and procrastination.

People are not always naive about this conflict. A sophisticated present-biased person predicts future temptation and seeks commitment. Alarm clocks placed across the room, automatic transfers, deadlines, restricted savings accounts and pre-booked sessions all alter the later choice set. The person gives up flexibility because flexibility will otherwise be used against the plan.

A randomised study in the Philippines offered some bank clients an account that restricted withdrawals until a chosen date or target had been reached. More than a quarter of those offered it accepted. The product was attractive because it removed an option. In standard consumer theory, extra options cannot make a person worse off if they may be ignored. Under self-control problems, the option to withdraw is itself a temptation and a restriction can have value.

Commitment creates a commercial fork. A firm can sell tools that align present action with earlier intention, or it can exploit the same inconsistency. Automatic renewal, delayed cancellation, introductory prices and minimum terms all trade on the gap between signing up and reviewing later. The contract may finance a lower price or useful service. It may also turn procrastination into revenue.

The welfare question depends on which self receives authority. The planner may be informed and calm, or merely ambitious. The acting self may be tempted, or responding to new information that made the old plan obsolete. A rigid commitment can protect a goal and trap a changed person. Behavioural policy therefore needs more than a preference for earlier intentions. It needs evidence that the commitment is chosen, proportionate and easy to revise when circumstances change in ways the earlier plan did not anticipate.

Time inconsistency explains why advice is often weaker than design. Telling a person to save more leaves the decisive action in the moment when spending is tempting. Moving the transfer to payday, setting it in advance or requiring an active opt-out changes who gets to choose and when. The best moment to make a difficult future choice is often before the future arrives.

6. Fairness Belongs Inside the Payoff

Take £10 and give it to one of two strangers. The first proposes a split. The second can accept, in which case both receive the proposed amounts, or reject, in which case both receive nothing. A person concerned only with personal money should accept any positive amount. One pound is more than zero.

In ultimatum-game experiments, responders often reject small offers, and proposers anticipate this by offering more. The responder pays to punish a division judged unfair. The proposer gives up money to avoid rejection, to comply with a norm, or both. The simple selfish-payoff model misses the bargain because the monetary amounts are not the whole payoff.

Economists responded by putting social concerns inside utility. A person may dislike disadvantageous inequality, dislike being far ahead, reward kind intentions, punish hostile ones or value conformity with an identity. These motives need not be noble. Envy, spite and retaliation are social preferences too. The change is analytical: another person's outcome and the perceived reason for it can affect one's own choice.

Fairness is not a single universal percentage. Cross-cultural experiments have found substantial variation in offers, rejection and cooperation across societies with different institutions and livelihoods. Market integration, group dependence, norms and the experiment's meaning all matter. A result from university students is evidence about those participants in that task, not a human constant waiting to be exported.

Context also changes whether fairness survives competition. In a bilateral bargain, a worker may refuse an insulting wage. In a crowded market, another worker may accept it, and competition can push the outcome towards the selfish benchmark even when no participant is selfish. Institutions aggregate motives in ways that can hide them. The fact that a market price looks impersonal does not prove the people inside the market are.

Firms encounter social preferences whenever a transaction carries an

entitlement. Customers may accept a price increase caused by higher costs and condemn an equal increase that appears to exploit a shortage. Workers respond to relative pay, recognition and procedural fairness as well as the wage level. A bonus can motivate when it signals reciprocity and alienate when it replaces a relationship with a crude price.

This creates limits on profit seeking that standard demand curves may overlook. A company can raise a price and lose trust beyond the immediate sales response. It can win a legal dispute and damage future cooperation. It can impose a target that increases measured output while teaching employees that every unpriced contribution is foolish. Social meaning becomes an economic input because it changes effort, loyalty, bargaining and enforcement.

The correction cuts both ways. Appeals to fairness can conceal self-interest, and stated concern for others may exceed costly action. Laboratory generosity can change when anonymity, repetition, selection or stakes change. Behavioural economics should not replace the selfish caricature with an altruistic one. It should model the conditions under which social motives enter and how institutions translate them into outcomes.

Self-interest remains one force among several. The useful question is not whether people are selfish or fair. It is what they treat as the payoff in this relationship, what norm defines a fair outcome, and whether the market makes that motive visible, profitable, punishable or irrelevant.

7. Choice Architecture Turns Behaviour Into an Economic Asset

The first Core Idea stripped rational choice down to a benchmark. The next five put the omitted person back in: scarce attention, reference points, frames, time conflict and social motives. The seventh is where those mechanisms become an economic system. Once behaviour depends on the environment, whoever controls the environment can create value, transfer value or extract it.

Every choice has an architecture even when nobody uses the word. A

pension scheme needs a rule for employees who do nothing. A subscription needs a renewal rule. A website must order its options. A restaurant chooses whether the expensive bottle appears first or last. A benefits agency decides how many documents an applicant must find. Neutrality is impossible because inaction, order, timing and friction always have consequences.

A nudge deliberately alters that architecture while preserving the main options and avoiding a large change in material incentives. Automatic enrolment, reminders, simplification and active-choice prompts fit the idea. Their attraction is obvious: if a small procedural change helps people carry out an intention they already hold, the designer can improve outcomes without a ban, subsidy or command.

The commercial version is broader than nudging. There are three distinct sources of behavioural profit. The first is **decision-cost reduction**. A firm can make search, comparison, remembering or execution easier. One-click reordering, sensible defaults, transparent bundles and useful reminders can be valuable because attention is scarce. The customer would still choose the service if the mechanism were fully explained.

The second is **commitment value**. A product can help a customer bind a future self to a present plan: automatic saving, prepaid sessions, spending limits, waiting periods or deadlines. The restriction can be part of what the customer buys. Extra flexibility is not always valuable when flexibility is the route through which a later self abandons an earlier goal.

The third is **error rent**. Revenue can depend on inattention, mistaken forecasts, hidden attributes or costly exit. Automatic renewal of an unwanted service, a fee revealed only after comparison, an obstructed cancellation path or a teaser price followed by customer inertia can convert predictable mistakes into margin. Sludge and digital dark patterns belong here. The same behavioural knowledge that removes friction from a public form can add friction exactly where a seller benefits from delay.

In practice, some error rents persist. Learning is weak when the choice is

rare, outcomes arrive years later or quality is hard to attribute. Mortgages, pensions, insurance and funerals do not provide rapid repetitions. Switching costs can protect a bad design. Complexity can be jointly profitable if consumers compare visible prices and ignore hidden ones, giving every seller an incentive to keep the hidden dimension obscure. Competition can therefore remove exploitation, leave it untouched or sharpen it.

The same discipline applies to public policy. Large-scale evidence has made the field less romantic about small interventions. In a complete set of 126 randomised trials conducted by two major US nudge units, covering about 23 million people, the average effect was 1.4 percentage points. Comparable published academic studies averaged 8.7 points. Differences in interventions mattered, but selective publication and low statistical power explained much of the gap. A one-point change can still matter across millions of people. It should not be sold as magic.

Later meta-analytic disputes sharpened the lesson. One influential 2022 synthesis reported a positive average effect across choice-architecture interventions; reanalyses argued that publication bias could remove that average. A 2025 second-order meta-analysis again found positive effects overall but substantial heterogeneity across nudge types and settings. There is no useful scalar called 'the nudge effect'. The unit of analysis must be the intervention, population, institution, outcome and time horizon.

Then comes the welfare problem created by the whole book. If observed choice partly reflects framing, inertia and self-control, which choice should the architect treat as the person's real interest? The earlier plan may be informed or merely aspirational. The active choice may express new information or temporary temptation. A government target can conflict with a citizen's aims; a seller's conversion metric can conflict with the customer's long-run welfare.

Core Idea 1 gave economics a clean person so choice could reveal preference. Behavioural economics showed why that inference can fail. Core Idea 7 is the bill for that discovery: if environments partly produce choices, the people who design environments acquire economic power. The

serious questions are therefore inseparable. What changes behaviour, does it work here, who profits, and whose objective counts as success?

How It Actually Works

Before economic man lost his nerves

Adam Smith's first great book was not about markets. *The Theory of Moral Sentiments*, published in 1759, examined sympathy, status, self-command, resentment and the human tendency to judge our own interests differently when they are near. Smith understood that people care about how they appear to others, overrate their prospects and feel some changes more sharply than equal changes in the opposite direction. The author later turned into a cardboard patron saint of selfish calculation had written an account full of social and psychological friction.

Early economics did not lack human observation. It lacked a compact way to turn all of it into tractable predictions. As the discipline became more mathematical in the twentieth century, it gained power by narrowing the person. Preferences, constraints and beliefs could be represented without explaining where desires came from or how attention worked. The omission was a design choice before it became a habit.

The useful simplification

Three moves built the modern benchmark. Revealed-preference theory treated consistent choices as evidence of underlying rankings. Expected-utility theory supplied rules for coherent choice under risk. Later work formalised beliefs and strategic interaction. Together they allowed economists to model a consumer, investor or firm without asking for introspection.

The achievement should not be minimised. A theory that predicts demand from prices and budgets is more useful than a catalogue of moods. Expected utility says, roughly, that a person facing lotteries behaves as though assigning utilities to outcomes, weighting them by probabilities and

choosing the highest total. Its axioms make choices internally coherent. They do not require the chooser to perform the arithmetic consciously.

The clean model also created a test. If people violated an axiom in a patterned way, the pattern might deserve a theory rather than dismissal as noise.

It also created a common language. A departure could be described against a precise prediction rather than against an observer's intuition. This mattered later, because behavioural economics would be strongest when it could say exactly which consistency condition failed and weakest when it merely announced that people were complicated. A surprise without a benchmark was a curiosity; a patterned violation could become economics. That distinction kept the new field answerable to prediction.

The cracks: certainty, ambiguity and good enough

Maurice Allais supplied one famous crack in 1953. He offered choices between lotteries arranged so that expected-utility theory required a consistent attitude across both pairs. Many people chose the safer option in one pair and the riskier option in the other. The attraction of certainty appeared to receive special weight. Allais was not showing that people chose badly. He was showing that a leading representation of rational risk choice could not describe the choices together.

Daniel Ellsberg attacked a different assumption in 1961 with an urn. Thirty balls were known to be red. The remaining sixty were black or yellow in an unknown proportion. People often preferred betting on the known chance of red to betting on an uncertain colour, then made another choice that could not be reconciled with one stable subjective probability. They treated measurable risk differently from ambiguity about the odds.

Herbert Simon had already widened the attack. In 1955 he argued that decision-makers cannot optimise over worlds they do not know and cannot compute. They search, use aspiration levels and stop at an option that meets requirements. Bounded rationality was not a theory of stupidity. It was a theory of decision costs inside real organisations. A manager with

incomplete data and a deadline does not solve an abstract optimum badly. She solves a different problem.

These challenges remained scattered. Allais concerned risky choice, Ellsberg uncertain beliefs, Simon limited search. Behavioural economics became a field when experiments began to reveal common mechanisms across many choices and when economists built those mechanisms into models of markets and policy.

Kahneman and Tversky change the question

Daniel Kahneman and Amos Tversky began collaborating in Jerusalem around 1969. Kahneman brought work on attention and perception. Tversky brought mathematical psychology and a taste for formal structure. Their method was unusually direct: construct a problem whose normative answer was clear, vary one feature, and see whether judgement moved in a way the benchmark forbade.

Their 1974 *Science* article organised judgement under uncertainty around representativeness, availability and anchoring. The importance was not the naming. Each heuristic connected a family of errors to an ordinary mental operation. People did not randomly neglect base rates. They answered a resemblance question. They did not randomly exaggerate dramatic risks. They sampled memory, whose contents were selected by vividness and recency. They did not randomly cluster estimates around a starting number. They adjusted from it.

Risky choice required another model. In 1979 Kahneman and Tversky published prospect theory. The paper separated an editing phase, in which options are framed and simplified, from evaluation. Outcomes became gains and losses from a reference point. The value function gave diminishing sensitivity and greater steepness for losses. Probabilities entered through decision weights rather than direct multiplication.

Prospect theory explained patterns expected utility handled poorly, including certainty effects, risk seeking over some losses and the importance of presentation. It was not a complete account of every risky choice. Its

achievement was to make context mathematically relevant without abandoning prediction.

Framing followed. In the 1981 disease problem, mathematically equivalent descriptions produced sharply different risk preferences. The result cut into the revealed-preference project. If a harmless change of wording reversed a choice, the economist could no longer assume that the action exposed a finished preference untouched by the elicitation method.

Thaler turns anomalies into economics

Richard Thaler's contribution was to ask what these findings did outside the psychology experiment. He collected economic choices that looked minor or embarrassing to the standard model: people who refused to sell an object for a price they would never pay to acquire it, households that treated one pound differently depending on its mental label, and consumers who ignored opportunity costs while reacting strongly to out-of-pocket losses.

Mental accounting supplied a framework for the informal budgets and transaction categories through which people controlled spending. The endowment effect connected ownership to reference dependence. Status quo bias and transaction utility showed how the terms of an exchange carried meaning beyond final consumption.

This was the moment behavioural economics became more than psychology applied to money. Thaler and others built models whose altered assumptions generated predictions about saving, asset trading, pricing and contracts. The aim was not to append a bias to every unexplained fact. It was to replace one assumption at a time and keep enough formal discipline to ask what followed.

The relationship with mainstream economics was adversarial but productive. A behavioural claim had to identify a regularity the standard model could not absorb cheaply, specify a mechanism and survive attempts to explain it through information, incentives or transaction costs. Some anomalies weakened under better procedures. Others became part of the discipline's ordinary toolkit.

Thaler's strategy was incremental. Keep optimisation where it worked, then alter the smallest assumption needed to explain the anomaly. If consumers separated money into accounts, model the accounts and ask how pricing changed. If investors reacted to gains and losses from a purchase price, specify the reference point and derive trading behaviour. This discipline distinguished behavioural economics from a loose appeal to human complexity. A model became better only if its added psychology earned predictive work.

The laboratory discovers the other person

A parallel experimental tradition tested markets and games. Vernon Smith showed that laboratory markets with real incentives could converge towards competitive predictions even when traders had limited information. This mattered because it blocked a crude conclusion. Individually imperfect people can generate orderly market outcomes. Institutions can discipline behaviour.

The ultimatum game, introduced experimentally in the early 1980s, revealed the reverse possibility. A responder could reject an unfair split and destroy both players' payoffs. Rejections of low offers showed that monetary self-interest did not describe every bilateral bargain. Models of inequity aversion, reciprocity and social preferences followed.

The two traditions belonged together. One showed that market rules can make the rational benchmark work better than the people appear to work alone. The other showed that social motives can survive incentives and change bargaining. Behavioural economics shifted attention from a permanent human type to the interaction between motives and institutions.

The pairing became official in 2002, when the economics Nobel was shared between Kahneman and Vernon Smith. One had shown systematic departures from the rational benchmark in individual judgement. The other had shown how experimental markets could produce orderly results. The award captured a tension that popular accounts often miss. Behavioural findings do not imply that markets are irrational crowds. Market rules,

feedback and competition can aggregate or discipline individual behaviour, while poorly designed markets can magnify the same tendencies.

Cross-cultural experiments deepened the lesson. Bargaining and cooperation varied across communities with different economic organisation and norms. The laboratory had revealed reproducible tasks, but not one fixed response shared at the same strength everywhere. Culture and institutions were not background noise. They helped determine what the game meant.

The field leaves the laboratory

The decisive expansion came from administrative records and field experiments. Instead of asking students to choose hypothetical gambles, researchers could observe pension enrolment, gym attendance, debt repayment, charitable giving and tax compliance.

Automatic enrolment in retirement plans supplied a clean institutional change. Under an opt-in rule, not acting meant not saving. Under an opt-out rule, not acting meant saving at the default rate and allocation. Participation rose, and many workers remained at the default. The same workers, incentives and options produced different outcomes because inertia had been reassigned.

The health-club study connected contract choice with later attendance and cancellation. It exposed overoptimism and delay through repeated behaviour rather than a single answer. The Philippine commitment-savings experiment went further by randomising access to a product that restricted future withdrawals. Take-up showed demand for self-control, and later balances showed that changing the choice set could change economic outcomes.

Behavioural finance, labour economics, development economics and public finance adopted related methods. Researchers studied limited attention to taxes, price endings, deadlines, overconfidence, reference-dependent labour supply and the response of firms to consumer mistakes. The field became less interested in proving that a bias existed and more interested in

estimating its size where money and institutions were real.

Financial markets offered an especially hard test. If an investor mispriced an asset, an informed trader might profit by correcting the price. Yet correction can be risky, costly or slow, and the mispricing can deepen before it closes. Work on overreaction, underreaction and excessive trading showed how beliefs and reference points might affect markets without assuming that every price movement was a psychological error. The challenge was identification: a pattern in returns can have many causes, and an attractive story about bias is not evidence on its own.

Firms also adapted. Some educated customers, simplified products or guaranteed outcomes because reducing error created value. Others designed contracts around overoptimism, inattention or cancellation delay. Competition did not choose one response automatically. If confused customers could compare easily, exploitation invited entry and reputation loss. If complexity hid the true price and every seller benefited from the confusion, competition could intensify the trick. The market response became part of the behavioural mechanism.

This also revealed selection. A gym sells contracts to people who choose gyms. A commitment account attracts people who recognise a self-control problem. A default affects people near the margin of action more than people with settled preferences. Average effects therefore combine mechanisms with the population exposed to them.

The nudge enters the state

The publication of *Nudge* by Thaler and Cass Sunstein in 2008 gave the field a public programme. If architecture cannot be avoided, they argued, it can be designed to help people make choices they would endorse while keeping alternatives open. Governments created behavioural-insights teams, and agencies began testing reminders, simplified letters, social comparisons and defaults.

The attraction was political as well as scientific. A redesigned form could improve take-up without a new benefit. A reminder could increase payment

without a larger fine. A default could change saving without a mandate. Randomised trials allowed rapid comparison at relatively low cost.

Commercial designers were running the mirror image. Online businesses could test button labels, order effects, scarcity messages, subscription flows and cancellation friction at enormous speed. The same experiment that found a helpful default could find a profitable trap. Behavioural knowledge entered a world where the architect often knew more about the chooser's pattern than the chooser knew about the experiment.

Consumer regulators responded with the language of sludge and dark commercial patterns. The issue was no longer whether a frame could influence choice. It was whether interface design impaired autonomy, concealed costs or converted inattention into revenue. Behavioural economics had moved from diagnosing anomalies to governing the institutions that manufacture them.

That move exposed a dispute about welfare. A rational-choice economist could often treat a voluntary purchase as evidence that the buyer expected to benefit. A behavioural economist could not do so automatically, because the purchase might depend on a hidden fee, a mistaken forecast or an obstructed exit. Yet rejecting revealed preference left no effortless replacement. Researchers proposed standards based on informed choice, stable plans, active decisions and the person's own considered judgement. Each standard helps in some cases and remains contestable in others.

By the time Thaler received the economics Nobel in 2017, the rebellion had entered the curriculum. Behavioural assumptions appeared in models of household finance, labour supply, industrial organisation and public economics. The shift was less a replacement of one theory than an expansion of the model-builder's menu. A standard assumption remained preferable when it predicted well and kept the analysis clear. A behavioural assumption earned inclusion when it explained a stable pattern, generated a new prediction or changed a policy conclusion. The field's success could be measured by how rarely an economist now needed to announce that attention, reference points or present bias were admissible topics.

The correction corrects itself

Success made the field vulnerable to overstatement. A memorable experiment was easier to publish than a null result. Small studies produced unstable estimates. Researchers had many analytical choices. Popular accounts mixed robust economic findings with fashionable psychological effects that later weakened.

Large replication projects changed the evidential standard. Coordinated teams repeated selected economics and social-science experiments with larger samples and common criteria. A majority of effects returned, but several did not, and successful replications were commonly smaller than the originals. Those results were neither vindication nor demolition. They showed a literature containing real effects, false positives and inflated estimates.

Complete portfolios from government nudge units made the same point operationally. When routine administrative trials, including null results, were analysed together, average effects were positive but much smaller than the academic literature had suggested. Better-powered administrative trials captured the modest interventions that journals had little reason to celebrate. Institutional context became part of the estimate rather than an afterthought.

The field's mature form is therefore less theatrical. It asks for preregistration, adequate power, repeated settings, heterogeneity, field outcomes and mechanisms that predict where an effect should change. It treats a named bias as the start of an investigation rather than the end.

Behavioural economics began by challenging an over-clean model of people. Its own progress has required the same treatment: fewer universal claims, more attention to selection and a sharper account of the institutions that turn a psychological tendency into an economic result.

How we know

The evidence comes from several methods that answer different questions. Controlled experiments can isolate a frame, default or payoff while holding other features fixed. Incentivised economic games reveal choices with real consequences, but the laboratory's rules, stakes and participants shape what the task means. Field experiments test interventions in working institutions, while administrative data connect choices to later behaviour at scale.

No method supplies a universal effect size. Student samples are narrow, field participants select into settings, and institutional changes often alter several mechanisms together. Published studies have favoured striking positive results, so preregistration, replication and complete trial registries matter. Cross-cultural work shows that social and bargaining behaviour can vary sharply across communities. Experience and competition can reduce some effects while leaving others intact.

The strongest claim is therefore conditional. Behavioural mechanisms are established when they recur under controlled tests, predict field behaviour and change in the expected direction when the environment changes. A finding that works once is an observation. A mechanism earns trust by explaining where it should work, where it should shrink and why.

What People Get Wrong

"Traditional economics says everyone is a perfect calculator"

The rational-agent model is a representation of consistent choice, not a claim that shoppers solve optimisation problems in the aisle. A person may use a habit, rule or adviser and still behave as though selecting the preferred feasible option. Thinking itself has a cost, so refusing to calculate a tiny saving can fit a wider rational model.

The misconception became persuasive because textbooks often introduced

the clean model before explaining its purpose, and behavioural writers gained drama by turning an abstraction into a ridiculous person. The correction matters because a behavioural explanation must beat a serious benchmark. If search costs, information, incentives or constraints explain the choice, adding a bias may make the model more colourful and less accurate. Behavioural economics is strongest when it identifies a specific failure of belief, attention or consistency, not when it mocks a calculation no economist claimed occurred. The standard model also remains the right tool in many settings. A correction earns its complexity by changing a prediction, a welfare judgement or the design of an institution. Otherwise the clean benchmark should remain as the default starting point for analysis in real institutions.

"A bias proves that people are irrational"

A bias is a systematic departure from a stated benchmark in a stated task. It is not a stable defect stamped on a person. The same shortcut can create error in one environment and perform well in another. Recognition, imitation and simple stopping rules often save time and can improve accuracy when data are noisy.

Some apparent biases also depend on misunderstanding the experiment, weak incentives or an unnatural representation of the problem. That does not make the findings worthless. It means the mechanism must survive changes designed to distinguish error from adaptation. The useful question is conditional: which cue is the person following, and does that cue track the outcome here? Calling the chooser irrational ends the investigation at the point where behavioural economics should begin it. The label is especially weak when applied after the outcome is known. Nearly any failure can be assigned a bias in retrospect. Prediction requires stating in advance which cue will dominate and how a changed environment should alter the error. That is the practical difference between a tested mechanism and a memorable nickname.

"Losses always hurt twice as much as gains"

The ratio has escaped from models into folklore. Prospect-theory estimates often use a value function steeper for losses than gains, and a coefficient around two became memorable. It was never a biological constant applying to every person, good, stake and reference point.

Loss aversion depends on what counts as a loss, whether ownership feels established, how choices are elicited, what experience the person has and whether transaction costs or misunderstanding are present. Some endowment-effect experiments produce large gaps; improved procedures or market experience can reduce them. The durable insight is that reference-dependent asymmetry often matters. The careless version replaces a conditional theory with a universal number and then uses the number to explain anything. A serious application must identify the reference point and estimate the response in the setting. It should also test rival explanations such as transaction costs, confusion and strategic bargaining. The phrase loss aversion is useful when it narrows the mechanism. It is useless when it is a flattering synonym for hindsight. The number should be estimated locally from relevant behaviour, not recited from memory.

"A default changes nothing except paperwork"

A default determines what happens when a person does not make an active choice. That makes it economically relevant whenever attention, effort, uncertainty or procrastination prevent action. Evidence from pension plans shows that moving enrolment into the path of inaction can alter participation sharply, even though workers retain the same formal right to refuse.

The mechanism is not always inertia. A default can look like expert advice, define the normal option, shift responsibility or reduce the fear of making a mistake. Those routes matter because they predict different remedies and risks. A well-chosen default can help people carry out an accepted goal. A poor default can lock passive people into unsuitable outcomes. Treating it as clerical detail ignores the cost of action; treating it as irresistible mind

control ignores active preferences and easy opt-outs. The practical test is to compare defaults, active-choice designs and later outcomes. Participation today may rise while contribution quality, persistence or subgroup welfare tells a more complicated story tomorrow. Defaults can move responsibility, expectations and action together, often without notice.

"Nudges are cheap substitutes for real policy"

A reminder cannot pay an unaffordable bill. A clearer menu cannot correct a monopoly price. A default cannot clean polluted air when emitting remains profitable. Behavioural design works on attention, friction, interpretation and action, so it fits problems where those mechanisms block an otherwise feasible choice.

At scale, nudge effects are often modest and heterogeneous. That can still justify them when implementation is cheap and millions of decisions are involved. The mistake is using a small design change to avoid a tax, standard, subsidy, service or enforcement action that addresses the cause. Good policy compares tools. A nudge may complement a financial incentive or simplify access to a right. It becomes political theatre when its light touch is the reason it was chosen despite weak fit. The relevant comparison is not nudge versus nothing. It is nudge versus the best available mix of prices, rules, information, capacity and enforcement. Behavioural insight should improve that comparison rather than pre-select the least demanding instrument. Low political cost is never evidence of high or durable social value.

"Businesses can exploit every bias forever"

Predictable mistakes can create profit, but a mistake is not automatically a business model. Exploitation survives only when the institution protects it.

Start with feedback. A customer who repeatedly buys the same simple product learns quickly whether it works and what it costs. A transparent rival can advertise the difference, comparison tools can expose hidden charges and reputation can punish obvious traps. In such markets, firms may profit more by reducing error than by harvesting it.

Now change the environment. Make the product complex, the choice infrequent, the outcome delayed and switching costly. Customers cannot cheaply learn from repetition. If everyone compares the headline price while neglecting an add-on, sellers can compete down the visible price and recover margin on the hidden one. Rivalry can preserve the distortion because the firm that becomes fully transparent may look expensive before the customer notices why.

The useful distinction is between value-creating behavioural profit and error rent. A good default, clear comparison or commitment feature can remain valuable after the customer understands it. Error rent depends on mistaken forecasts, hidden terms, inattention or obstructed exit. It may be profitable, but its durability depends on feedback, competition, reputation, regulation and the cost of switching. Behavioural economics predicts none of these away. It tells you to model them.

"The replication crisis destroyed behavioural economics"

The replication crisis damaged the field's confidence, not its subject matter. Some celebrated findings failed to reproduce, and repeated effects were often smaller than the originals. In one replication project covering eighteen laboratory economics studies, about two-thirds reproduced by the main criterion. A second project on twenty-one high-profile social-science experiments found significant effects in the same direction for thirteen, with replication effect sizes averaging about half the originals.

Nudging became a useful stress test. Large operational trials produced much smaller average effects than the published academic literature. A 2022 meta-analysis reported a positive average effect across choice-architecture interventions, while contemporaneous reanalyses argued that publication-bias adjustment could reduce the pooled estimate to zero. A 2025 second-order meta-analysis found positive effects overall but strong heterogeneity by intervention and setting. Social-norm messaging in health provides an even narrower warning: a 2025 preregistered meta-analysis found a small unadjusted effect that disappeared after

publication-bias adjustment.

Those findings do not support a field-wide pass rate or obituary. Behavioural economics contains robust mechanisms, procedure-sensitive findings, heterogeneous effects and claims that should be abandoned. The methodological response is more important than the score: preregistration, larger samples, complete trial sets, independent repetition, field evidence and explicit theories of when an effect should change. The mature question is no longer 'is this bias real?' but 'what mechanism survives which test, at what size, in which environment?'

Use It

Find the benchmark before naming the bias

When a choice looks foolish, resist the satisfaction of diagnosis. State the person's likely objective, information, constraints and available alternatives. Then describe what a coherent choice would require. Only after that should you name the departure.

This prevents three common errors. A liquidity constraint can look like impatience. A high search cost can look like indifference to price. A preference you dislike can look like a cognitive failure. The benchmark may also expose a real puzzle: a customer pays extra for easy cancellation, repeatedly fails to use the service and then delays exit. Constraints alone do not explain why the chosen flexibility remains unused.

In business, ask what a customer would choose if prices were transparent, comparison were easy and attention were available. In policy, ask whether non-take-up reflects rejection, ignorance or process cost. A bias label is useful only when it changes the predicted response to a redesign. Write that prediction down before changing the process. Otherwise improvement and storytelling become impossible to separate.

Locate the reference point

Before judging a reaction to money, risk or performance, ask what outcome counts as zero. It may be last year's pay, a promised delivery date, a purchase price, a target, a competitor's result or the best outcome recently imagined.

Many conflicts that look disproportionate become intelligible once the reference point appears. A customer receiving a refund can still feel cheated if full performance had become the entitlement. An employee can resent a pay rise below expectation. A founder can keep funding a failing project because closure would convert a paper disappointment into a recognised loss.

Reference points can be influenced, but doing so carelessly creates later costs. A discount trains customers to treat the lower price as normal. An aggressive target may motivate until missing it becomes a loss that invites distortion. Good analysis records how the reference point was formed and how quickly it may adapt.

Separate preference from friction

A completed action mixes desire with the cost of acting. An uncompleted action does too. The fact that someone stayed subscribed does not prove satisfaction if cancellation requires a phone call, repeated screens and a final-day reminder. The fact that someone failed to claim a benefit does not prove the benefit was unwanted if the form required documents, time and confidence.

Run the reversal test. Make the opposite option the default, equalise the number of steps, provide the same information and see whether behaviour changes. Large movement suggests that the original action contained friction or implied advice as well as preference.

This lens is useful beyond interfaces. Meeting attendance, staff surveys, pension choices, renewals and complaints all depend on process design. Measure the path, not merely the endpoint. Friction is sometimes necessary

for security or deliberation, but it should be placed where the designer can defend it rather than where it quietly protects revenue or administrative convenience.

Make the future decision early

When a plan repeatedly loses at the moment of temptation, stop demanding stronger character from that moment. Move the decision.

Set transfers for payday rather than asking the end-of-month self to save what remains. Book the session before the tired evening arrives. Use cooling-off periods for purchases made under urgency and commitment devices for goals chosen under reflection. The aim is to give the preferred self a procedural advantage.

Do not confuse commitment with punishment. A useful device is voluntarily chosen, proportionate and revisable when circumstances change. It removes the temptation that defeats a stable plan without making ordinary mistakes ruinous. The strongest commitment may be social, financial, technological or environmental. Its design should match the failure mechanism: forgetting needs a reminder, procrastination needs a deadline or default, and temptation may need distance or restriction.

Follow the architect's incentive

Every choice environment has an owner. Identify what that owner is paid, measured or blamed for. The answer often predicts the direction of the design better than the language around it.

A pension trustee may prefer adequate long-term saving. A retailer may prefer conversion today. A public agency may prefer completed applications, lower cost or lower recorded demand. Those goals can align with the chooser's interests and can conflict. A large green button, a pre-ticked box or a difficult exit has no moral quality by itself; its function depends on the objective it advances and the information it conceals.

Inspect asymmetry. Is joining easier than leaving? Are gains prominent and costs deferred? Does the designer collect the upside while the chooser

bears mistakes? Transparent, reversible architecture is harder to abuse because the mechanism remains visible and the person can correct it.

For a business, distinguish value-creating behavioural profit from error rent. Value-creating profit reduces search, uncertainty or self-control costs and can survive when the customer understands the mechanism. Error rent depends on mistaken forecasts, hidden terms or obstructed exit. It may lift conversion while increasing complaints, regulation and future acquisition costs. The source of profit predicts its durability.

Demand local evidence

Do not ask whether nudges work. Ask whether this intervention changes this behaviour for this population under this institution, and what happens afterwards.

A reminder may help people who intended to act and forgot. It may irritate people who cannot act. A default can lift participation while producing unsuitable contribution levels. An effect measured at sign-up can disappear at retention. A result from students choosing hypothetical gambles may identify a mechanism without estimating its commercial size.

Prefer complete trial sets, adequate samples, preregistered outcomes and measures tied to welfare rather than clicks. Look for heterogeneity: an average gain can conceal harm to a subgroup or no effect on the people the intervention was meant to reach. Repetition across settings matters, but variation is information. A mechanism earns confidence when its effect changes for reasons the theory predicted.

The limits

Behavioural economics cannot tell you what people should want. It can identify inconsistency, inattention and manipulation, but many disagreements are about values rather than errors. A person may knowingly prefer convenience, risk, loyalty or pleasure over the outcome an analyst calls efficient.

The field can also become unfalsifiable. If every result receives a bias after

the event, the vocabulary explains nothing. Reference points can be selected to fit the outcome. Multiple mechanisms can predict the same action. Researchers and practitioners are subject to confirmation bias, selective attention and incentives to report success.

Do not use behavioural language to shrink a structural problem into a personal one. Poverty is not a failure to budget. Monopoly is not a framing effect. Unsafe work is not a shortage of resilience. Attention and choice matter inside institutions whose prices, rights and power may dominate the result.

Finally, knowledge does not confer immunity. The person who can define anchoring is still exposed to anchors. Expertise helps most when it changes the environment, the process or the evidence, not when it produces confidence that the mechanism applies only to other people.

The one thing to keep

Keep the economic architecture in view.

A decision feels private because the last movement occurs inside one person. Yet the available options, the comparison price, the default, the deadline, the information order and the cost of reversal were usually built elsewhere. The chooser supplies motives, habits and limits. The institution decides which of them are cheap to express.

That does not make people puppets. They learn, resist, search and hold preferences that designers cannot wish away. It does mean that a transaction is never explained by the chooser alone. A subscription may reveal genuine demand and procrastination. A pension contribution may reveal a wish to save and a default. A low claim rate may reveal weak demand and an exhausting form. The observed choice is an equilibrium between person and environment.

So when money, policy or welfare matters, ask four questions. What would the rational benchmark predict? Which behavioural mechanism could move the choice? What does the designer gain if behaviour moves that way?

What evidence would distinguish value creation from error extraction?

That last distinction is the subtitle's permanent lesson. Behavioural economics can make you money, but the mechanism matters. Profit earned by reducing decision costs or helping people carry out their own plans can survive transparency. Profit earned from hidden terms, mistaken forecasts or obstructed exit depends on the conditions that keep the mistake alive.

The field's deepest achievement is therefore not a catalogue of irrationality. It is a more demanding model of choice. People decide with limited attention, shifting reference points, competing selves and social motives inside environments designed by other people with incentives of their own. Better economics begins when both sides of that sentence enter the model.

Terms

Homo economicus. The simplified economic actor with coherent preferences and beliefs who chooses the best feasible option. It is a modelling device, not a claim about calculation, selfishness or moral worth.

Rational choice. Choice that is internally consistent with stated preferences, beliefs and constraints. The standard varies by model, so calling behaviour irrational requires specifying which condition has failed.

Revealed preference. The inference that choices expose an underlying ranking of options. Behavioural economics complicates it by showing that frames, mistakes and friction can alter what the same person chooses.

Expected utility. A model of risky choice in which utilities of possible outcomes are weighted by their probabilities. Its axioms provide a powerful benchmark that several famous experiments systematically violate.

Bounded rationality. Herbert Simon's account of decision-making under limited information, time and computation. It directs attention from impossible optimisation towards search rules, aspiration levels and organisational constraints.

Satisficing. Stopping a search when an option meets an acceptable

threshold rather than proving it is best. The rule can be sensible when further search costs more than likely improvement.

Heuristic. A decision rule that ignores some information or substitutes an easier judgement. Heuristics save effort and can be accurate when their cues fit the environment.

Availability heuristic. Estimating frequency or probability from how easily examples come to mind. Vividness, repetition and recency can therefore move perceived risk without changing the underlying rate.

Representativeness heuristic. Judging probability by resemblance to a type or story. It becomes misleading when similarity overwhelms base rates, sample size or the process that generated the evidence.

Anchoring. Starting from an available number and adjusting insufficiently. Anchors can be informative, arbitrary or strategically supplied, which makes first offers and displayed prices economically consequential.

Base rate. The underlying prevalence of an event or category before case-specific evidence is considered. Ignoring it can make a persuasive match look more diagnostic than it is.

Prospect theory. Kahneman and Tversky's model of risky choice using reference-dependent gains and losses, diminishing sensitivity and decision weights. It was built to explain patterns expected utility missed.

Reference point. The outcome treated as the comparison baseline, such as the status quo, expectation, target or purchase price. Moving it can convert the same result from gain to loss.

Loss aversion. Greater sensitivity to losses than equal gains around a reference point. Its size varies by setting and should not be treated as a universal two-to-one law.

Diminishing sensitivity. The tendency for equal changes to matter less when they occur farther from the reference point. The first £100 commonly changes value more than another £100 added to thousands.

Probability weighting. The transformation of objective probabilities into decision weights. Small chances may receive disproportionate attention, while moderate and high probabilities are often compressed rather than treated linearly.

Framing effect. A change in choice caused by how equivalent options are described, organised or compared. It matters because rational preference should be invariant to irrelevant presentation.

Endowment effect. Higher valuation associated with possessing or feeling entitled to an object. The observed gap depends on reference points, procedures, experience and what the transaction means.

Status quo bias. A tendency to retain the current option beyond what switching costs alone predict. Inertia, loss aversion, implied endorsement and responsibility can all contribute.

Default. The outcome applied when no active choice is made. Defaults allocate the consequences of inaction and can communicate advice, reduce effort or protect the designer's preferred result.

Mental accounting. Informal bookkeeping that labels money and separates transactions into accounts. It helps control behaviour but violates the economic principle that equivalent pounds are interchangeable.

Present bias. Extra weight placed on immediate costs and rewards relative to later ones. It can reverse earlier plans when the moment of action arrives.

Hyperbolic discounting. A family of models in which discounting falls sharply near the present and more gently later. It generates time inconsistency and demand for commitment.

Commitment device. A chosen restriction, cost or rule that protects a plan from later temptation or procrastination. Its value comes from removing flexibility that the future self may misuse.

Social preference. A preference in which another person's outcome, intentions or the fairness of a distribution affects utility. Cooperation, envy,

reciprocity and punishment can all fit the category.

Reciprocity. Responding to perceived kindness with kindness and hostility with punishment, sometimes at personal cost. Intentions and entitlements therefore influence exchange beyond final monetary payoffs.

Ultimatum game. A bargaining experiment in which a responder may reject a proposed split and leave both players with nothing. Rejection of low offers reveals costly concern with fairness or intent.

Choice architecture. The design of the environment in which options are presented and acted upon. Order, defaults, timing and friction can change behaviour without changing the formal menu.

Nudge. A choice-architecture intervention that predictably changes behaviour while preserving options and avoiding major incentive changes. Its effect is specific to an intervention, population and setting.

Sludge. Friction that makes a beneficial or intended action harder, such as claiming, comparing, cancelling or complaining. Sludge may arise from neglect or be designed to protect revenue or administrative convenience, and its burden often falls unevenly across time, income, disability and confidence.

Go Deeper

The accessible history

Richard H. Thaler, *Misbehaving: The Making of Behavioral Economics* (2015). Thaler tells the field's development through the anomalies, arguments and collaborations that moved psychology into economics. It is funny, concrete and unusually clear about how a discipline changes. The viewpoint is that of a central participant, so it gives less space to critics and parallel traditions than a detached history would. Begin here if this book has made you curious rather than technical. Pay particular attention to the failed starts and hostile seminars: they show that evidence changes a field through argument and institutional persistence, not through one decisive experiment.

The primary theory

Daniel Kahneman and Amos Tversky, "Prospect Theory: An Analysis of Decision under Risk" (1979). This is the paper that made reference points, loss aversion and probability weighting into a formal alternative to expected utility. The examples remain readable, but the notation and model require slow attention. Read it to see how behavioural economics earns the final word in its name: a psychological regularity becomes valuable when it is stated precisely enough to generate economic predictions. The paper also reveals the model's limits, especially the work required to identify a reference point outside a controlled problem.

The evidence outside the laboratory

Stefano DellaVigna, "Psychology and Economics: Evidence from the Field" (2009). This long survey organises the subject into nonstandard preferences, beliefs and decision-making, then tests them against behaviour in consumption, finance, work, crime, voting and charity. It is the best bridge from famous experiments to economically consequential evidence. The article is scholarly but structured well enough for a patient general reader. Use its organisation as a map: preferences, beliefs and decision procedures are different sources of departure, and each produces different evidence and policy questions.

The strongest corrective

Gerd Gigerenzer and Wolfgang Gaissmaier, "Heuristic Decision Making" (2011). Read this after the biases literature. It explains why rules that ignore information can be efficient and, in the right environment, more accurate than heavier calculation. The paper does not erase systematic error. It prevents the lazy inference that every shortcut is a defect and directs attention towards the fit between a decision rule and its environment. It is the best antidote to reading a catalogue of biases as a catalogue of permanent human flaws. Read the examples with one question in mind: what structure in the environment makes this rule intelligent, and when does that structure disappear?

Notes and Sources

Compressed argument

The distinction between rationality as a benchmark and a literal description follows the standard use of revealed preference and expected utility, alongside Herbert Simon's bounded-rationality critique. The main behavioural families are drawn from Tversky and Kahneman's work on heuristics, prospect theory and framing; Thaler's work on mental accounting, ownership and choice architecture; research on present bias and social preferences; and DellaVigna's synthesis of field evidence.

The claim that effects vary by setting and scale is supported by replication projects led by Colin Camerer and by DellaVigna and Linos's complete set of nudge-unit trials. The commercial discussion uses the OECD's 2022 report on dark commercial patterns and the US Federal Trade Commission's report from the same year.

Relevance and opening evidence

The health-club figures come from DellaVigna and Malmendier's analysis of 7,752 members at three US clubs. Flat-fee members paid more than \$70 a month, attended 4.3 times a month on average and paid more than \$17 per expected visit against a \$10 pass. The reported average forgone saving was \$600. Monthly members were 17 per cent more likely than annual members to remain beyond a year despite paying for easier cancellation.

The pension-default claim is based on Madrian and Shea's study of automatic enrolment in a 401(k) plan. The authors found higher participation and substantial persistence at the default contribution and allocation. They discuss inertia and the interpretation of the default as advice.

Core-idea evidence

The rational benchmark. Samuelson's revealed-preference programme, von Neumann and Morgenstern's expected-utility framework and Simon's 1955 model provide the foundation. Simon used bounded rationality to redirect analysis towards the procedures people and organisations use when optimisation is infeasible. The point that rationality does not imply selfishness follows directly from treating social outcomes as arguments in a utility function.

Heuristics. Tversky and Kahneman's 1974 *Science* article describes representativeness, availability and anchoring. Gigerenzer and Gaissmaier review evidence that simple heuristics can exploit environmental structure and sometimes outperform more information-intensive methods. The manuscript therefore treats a heuristic as a rule whose quality depends on fit, feedback and the decision environment.

Prospect theory. Kahneman and Tversky's 1979 paper supplies reference dependence, diminishing sensitivity, different treatment of gains and losses, and transformed probability weights. The warning against a universal two-to-one loss ratio reflects later debate over parameter variation, reference-point identification and experimental procedure. Plott and Zeiler show that parts of the willingness-to-pay and willingness-to-accept gap are sensitive to elicitation procedures and subject understanding.

Frames, accounts and ownership. The disease problem comes from Tversky and Kahneman's 1981 framing article. Thaler developed mental accounting in work first published in the 1980s. Kahneman, Knetsch and Thaler's 1990 experiments are the main source for the mug-trading example, with Plott and Zeiler providing an important procedural challenge. The default evidence again comes from Madrian and Shea.

Time inconsistency. Laibson's 1997 model formalises hyperbolic discounting and demand for commitment. The gym evidence comes from DellaVigna and Malmendier. Ashraf, Karlan and Yin designed the Philippine commitment-savings product and tested it randomly. Of 710 clients offered the account, 202 accepted, and treatment assignment produced substantially higher later savings balances than the control condition.

Social preferences. Güth, Schmittberger and Schwarze introduced the experimental ultimatum-game evidence used here. Fehr and Schmidt provide a formal model of inequity aversion. Kahneman, Knetsch and Thaler supply the evidence on fairness, reference transactions and acceptable price or wage changes. Henrich and colleagues report substantial cross-cultural variation in bargaining and cooperation across small-scale societies, warning against treating one student sample as a universal response.

Choice architecture. Thaler and Sunstein developed the nudge programme for a wide audience. DellaVigna and Linos assembled 126 randomised trials covering 23 million people from two US nudge units. They estimate an average take-up effect of 8.7 percentage points in the academic comparison sample and 1.4 points in the nudge-unit sample. Differences in

power, intervention characteristics and selective publication explain much of the gap. Mertens and colleagues reported a positive meta-analytic average for choice architecture interventions; Maier and colleagues argued that adjusting for publication bias removed that average. A 2025 second-order meta-analysis by Hu and colleagues synthesised 14 earlier meta-analyses covering 1,638 primary studies and about 30 million participants. It found a small aggregate effect and an estimate near zero after publication-bias adjustment, while rating most contributing meta-analyses low or critically low quality. The evidence therefore reinforces the manuscript's insistence on intervention-specific testing rather than supplying one universal nudge effect.

Historical and methodological spine

Ashraf, Camerer and Loewenstein document the behavioural themes in Adam Smith's *Theory of Moral Sentiments* and correct the image of an intellectually psychology-free classical economics. The formal benchmark is represented by Samuelson, von Neumann and Morgenstern. Allais's 1953 paper and Ellsberg's 1961 paper establish the certainty and ambiguity challenges. Simon's 1955 article supplies bounded rationality.

Kahneman's Nobel lecture and the original papers provide the account of the collaboration with Tversky, the 1974 heuristics programme, prospect theory and framing. Thaler's *Misbehaving* is used for the development of mental accounting and the institutional history of the field, checked against the original articles. The 2002 Nobel award was shared between Kahneman for integrating psychological research into economic science and Vernon Smith for establishing laboratory experiments as a tool in empirical economic analysis. Thaler received the 2017 award for contributions to behavioural economics.

The external-validity discussion uses Levitt and List's analysis of what laboratory social-preference experiments reveal about the field, Henrich and colleagues' cross-cultural evidence, and DellaVigna's survey of field findings. The welfare discussion draws on Bernheim and Rangel's attempt to build behavioural welfare analysis beyond unqualified revealed

preference.

The replication figures are from Camerer and colleagues. Their 2016 economics project replicated roughly two-thirds of eighteen selected laboratory findings by the principal criterion. Their 2018 social-science project found a significant effect in the same direction for thirteen of twenty-one studies, with replication effects averaging about half of the originals. These projects do not sample every behavioural-economics claim and are used as evidence about research practice, not as a field-wide pass rate.

Misconception corrections

The first correction follows the distinction between as-if maximisation and conscious calculation. The second uses the ecological-rationality literature and experimental-procedure critiques. The loss-aversion correction reflects the original prospect-theory model, later reference-dependent work and challenges to universal interpretation of endowment experiments.

The defaults correction rests on Madrian and Shea and on the wider default literature. The policy correction uses DellaVigna and Linos's evidence that effects can remain useful at scale while being smaller than published estimates. The market correction follows DellaVigna's review of firm responses and the standard argument that competition, feedback and experience can discipline some errors but can also reward complexity. The replication correction uses the two Camerer projects and the nudge meta-analysis debate.

Applied lenses

The lenses translate the book's mechanisms rather than claim therapeutic or financial guarantees. The benchmark test follows standard economic identification. Reference-point analysis follows prospect theory.

Preference-friction separation follows defaults, administrative-burden research and consumer-contract evidence. Early commitment follows Laibson and the field experiments on gyms and savings. Architect incentives follow the nudge, sludge and dark-pattern literature. The demand for local evidence follows the field's replication, scale and heterogeneity findings.

Glossary

Definitions follow the works cited above. Nudge is used in its conventional behavioural-policy sense: an alteration of choice architecture that predictably changes behaviour without forbidding options or materially changing economic incentives. Sludge refers to process friction that obstructs action and may be intentional or accidental.

Recommendations

Publication details for all four recommendations were checked against the publishers or journal records. Thaler's book is used as an insider history rather than a neutral survey. Prospect theory is the primary technical work. DellaVigna provides the field-evidence synthesis. Gigerenzer and Gaissmaier provide the principal corrective to a bias-only account of heuristics.

Bibliography

Primary and original research

Allais, Maurice. "Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'école américaine." *Econometrica* 21, no. 4 (1953): 503-546. DOI 10.2307/1907921.

Ashraf, Nava, Dean Karlan, and Wesley Yin. "Tying Odysseus to the Mast:

Evidence From a Commitment Savings Product in the Philippines." *Quarterly Journal of Economics* 121, no. 2 (2006): 635-672. DOI 10.1162/qjec.2006.121.2.635.

Camerer, Colin F., et al. "Evaluating Replicability of Laboratory Experiments in Economics." *Science* 351, no. 6280 (2016): 1433-1436. DOI 10.1126/science.aaf0918.

Camerer, Colin F., et al. "Evaluating the Replicability of Social Science Experiments in Nature and Science Between 2010 and 2015." *Nature Human Behaviour* 2 (2018): 637-644. DOI 10.1038/s41562-018-0399-z.

DellaVigna, Stefano, and Elizabeth Linos. "RCTs to Scale: Comprehensive Evidence From Two Nudge Units." *Econometrica* 90, no. 1 (2022): 81-116. DOI 10.3982/ECTA18709.

DellaVigna, Stefano, and Ulrike Malmendier. "Paying Not to Go to the Gym." *American Economic Review* 96, no. 3 (2006): 694-719. DOI 10.1257/aer.96.3.694.

Ellsberg, Daniel. "Risk, Ambiguity, and the Savage Axioms." *Quarterly Journal of Economics* 75, no. 4 (1961): 643-669. DOI 10.2307/1884324.

Güth, Werner, Rolf Schmittberger, and Bernd Schwarze. "An Experimental Analysis of Ultimatum Bargaining." *Journal of Economic Behavior & Organization* 3, no. 4 (1982): 367-388. DOI 10.1016/0167-2681(82)90011-7.

Henrich, Joseph, et al. "'Economic Man' in Cross-Cultural Perspective: Behavioral Experiments in 15 Small-Scale Societies." *Behavioral and Brain Sciences* 28, no. 6 (2005): 795-815. DOI 10.1017/S0140525X05000142.

Kahneman, Daniel, Jack L. Knetsch, and Richard H. Thaler. "Experimental Tests of the Endowment Effect and the Coase Theorem." *Journal of Political Economy* 98, no. 6 (1990): 1325-1348. DOI 10.1086/261737.

Kahneman, Daniel. "Maps of Bounded Rationality: Psychology for Behavioral Economics." *American Economic Review* 93, no. 5 (2003): 1449-1475. DOI 10.1257/000282803322655392.

Kahneman, Daniel, Jack L. Knetsch, and Richard H. Thaler. "Fairness as a Constraint on Profit Seeking: Entitlements in the Market." *American Economic Review* 76, no. 4 (1986): 728-741.

Kahneman, Daniel, and Amos Tversky. "Prospect Theory: An Analysis of Decision under Risk." *Econometrica* 47, no. 2 (1979): 263-291. DOI 10.2307/1914185.

Laibson, David. "Golden Eggs and Hyperbolic Discounting." *Quarterly Journal of Economics* 112, no. 2 (1997): 443-478. DOI 10.1162/003355397555253.

Madrian, Brigitte C., and Dennis F. Shea. "The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior." *Quarterly Journal of Economics* 116, no. 4 (2001): 1149-1187. DOI 10.1162/003355301753265543.

Plott, Charles R., and Kathryn Zeiler. "The Willingness to Pay-Willingness to Accept Gap, the 'Endowment Effect,' Subject Misconceptions, and Experimental Procedures for Eliciting Valuations." *American Economic Review* 95, no. 3 (2005): 530-545. DOI 10.1257/0002828054201387.

Samuelson, Paul A. "A Note on the Pure Theory of Consumer's Behaviour." *Economica* 5, no. 17 (1938): 61-71.

Simon, Herbert A. "A Behavioral Model of Rational Choice." *Quarterly Journal of Economics* 69, no. 1 (1955): 99-118. DOI 10.2307/1884852.

Smith, Vernon L. "An Experimental Study of Competitive Market Behavior." *Journal of Political Economy* 70, no. 2 (1962): 111-137. DOI 10.1086/258609.

Thaler, Richard H. "Mental Accounting and Consumer Choice." *Marketing Science* 4, no. 3 (1985): 199-214. DOI 10.1287/mksc.4.3.199.

Tversky, Amos, and Daniel Kahneman. "Judgment under Uncertainty: Heuristics and Biases." *Science* 185, no. 4157 (1974): 1124-1131. DOI 10.1126/science.185.4157.1124.

Tversky, Amos, and Daniel Kahneman. "The Framing of Decisions and the Psychology of Choice." *Science* 211, no. 4481 (1981): 453-458. DOI 10.1126/science.7455683.

Modern works and syntheses

Ashraf, Nava, Colin F. Camerer, and George Loewenstein. "Adam Smith, Behavioral Economist." *Journal of Economic Perspectives* 19, no. 3 (2005): 131-145. DOI 10.1257/089533005774357897.

Bernheim, B. Douglas, and Antonio Rangel. "Beyond Revealed Preference: Choice-Theoretic Foundations for Behavioral Welfare Economics." *Quarterly Journal of Economics* 124, no. 1 (2009): 51-104. DOI 10.1162/qjec.2009.124.1.51.

DellaVigna, Stefano. "Psychology and Economics: Evidence from the Field." *Journal of Economic Literature* 47, no. 2 (2009): 315-372. DOI 10.1257/jel.47.2.315.

Fehr, Ernst, and Klaus M. Schmidt. "A Theory of Fairness, Competition, and Cooperation." *Quarterly Journal of Economics* 114, no. 3 (1999): 817-868. DOI 10.1162/003355399556151.

Federal Trade Commission. *Bringing Dark Patterns to Light*. Washington, DC: Federal Trade Commission, 2022.

Gigerenzer, Gerd, and Wolfgang Gaissmaier. "Heuristic Decision Making." *Annual Review of Psychology* 62 (2011): 451-482. DOI 10.1146/annurev-psych-120709-145346.

Hu, Bo, Ziqian Xia, Qingcheng Guo, Chang Lu, Sara M. Constantino, and Xingda Ju. "Assessing Nudge Impact: A Comprehensive Second-Order Meta-Analysis." *Journal of Behavioral Decision Making* 38, no. 5 (2025): e70053. DOI 10.1002/bdm.70053.

Papakonstantinou, Trisevgeni, Sarah Lynn Flecke, C. E. R. Edmunds, Rosina Cross, Anh Tran, et al. "A Systematic Review and Meta-Analysis of the Effectiveness of Social Norms Messaging Approaches for Improving

Health Behaviours in Developed Countries." *Nature Human Behaviour* 9 (2025): 2632-2650. DOI 10.1038/s41562-025-02275-6.

Levitt, Steven D., and John A. List. "What Do Laboratory Experiments Measuring Social Preferences Reveal About the Real World?" *Journal of Economic Perspectives* 21, no. 2 (2007): 153-174. DOI 10.1257/jep.21.2.153.

Maier, Maximilian, et al. "No Evidence for Nudging After Adjusting for Publication Bias." *Proceedings of the National Academy of Sciences* 119, no. 31 (2022): e2200300119. DOI 10.1073/pnas.2200300119.

Mertens, Stephanie, et al. "The Effectiveness of Nudging: A Meta-Analysis of Choice Architecture Interventions Across Behavioral Domains." *Proceedings of the National Academy of Sciences* 119, no. 1 (2022): e2107346118. DOI 10.1073/pnas.2107346118.

Organisation for Economic Co-operation and Development. *Dark Commercial Patterns*. OECD Digital Economy Papers, no. 336. Paris: OECD Publishing, 2022. DOI 10.1787/44f5e846-en.

Thaler, Richard H. *Misbehaving: The Making of Behavioral Economics*. New York: W. W. Norton, 2015.

Thaler, Richard H., and Cass R. Sunstein. *Nudge: The Final Edition*. New Haven: Yale University Press, 2021.

von Neumann, John, and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton: Princeton University Press, 1944.